

Conditional path sampling for stochastic differential equations through drift relaxation

Panos Stinis
 Department of Mathematics
 University of Minnesota
 Minneapolis, MN 55455

Abstract

We present an algorithm for the efficient sampling of conditional paths of stochastic differential equations (SDEs). While unconditional path sampling of SDEs is straightforward, albeit expensive for high dimensional systems of SDEs, conditional path sampling can be difficult even for low dimensional systems. This is because we need to produce sample paths of the SDE which respect both the dynamics of the SDE and the initial and endpoint conditions. The dynamics of a SDE are governed by the deterministic term (drift) and the stochastic term (noise). Instead of producing conditional paths directly from the original SDE, one can consider a sequence of SDEs with modified drifts. The modified drifts should be chosen so that it is easier to produce sample paths which satisfy the initial and endpoint conditions. Also, the sequence of modified drifts converges to the drift of the original SDE. We construct a simple Markov Chain Monte Carlo (MCMC) algorithm which samples, in sequence, conditional paths from the modified SDEs, by taking the last sampled path at each level of the sequence as an initial condition for the sampling at the next level in the sequence. The algorithm can be thought of as a stochastic analog of deterministic homotopy methods for solving nonlinear algebraic equations or as a SDE generalization of simulated annealing. The algorithm is particularly suited for filtering/smoothing applications. We show how it can be used to improve the performance of particle filters. Numerical results for filtering of a stochastic differential equation are included.

Introduction

The study of systems arising in different areas, from signal processing and chemical kinetics to econometrics and finance (see e.g. [1, 18]) often requires

the sampling of paths of stochastic differential equations (SDEs) subject to initial and endpoint conditions. While unconditional path sampling of SDEs is straightforward, albeit expensive for high dimensional systems of SDEs, conditional path sampling can be difficult even for low dimensional systems. This is because we need to produce sample paths of the SDE which respect both the dynamics of the SDE and the initial and endpoint conditions. An analogous situation arises in ordinary differential equations, where it can be considerably more difficult to create solutions to boundary value problems than it is to construct solutions to initial value problems (see e.g. Ch. 8 in [4]). The problem of conditional path sampling of SDEs has been a subject of active research in recent years and some very interesting approaches have already been developed (see e.g. [2, 18, 21]).

The dynamics of a SDE are governed by the deterministic term (drift) and the stochastic term (noise). Instead of producing conditional paths directly from the original SDE, one can consider a sequence of SDEs with modified drifts. The modified drifts should be chosen so that it is easier to produce sample paths which satisfy the initial and endpoint conditions. Also, the sequence of modified drifts converges to the drift of the original SDE. We construct a simple Markov Chain Monte Carlo (MCMC) algorithm which samples, in sequence, conditional paths from the modified SDEs, by taking the last sampled path at each level of the sequence as an initial condition for the sampling at the next level in the sequence.

We have used the drift relaxation algorithm to modify a popular filtering method called particle filter [5]. A particle filter is a sequential importance sampling algorithm which is based on the recursive (online) Bayesian updating of the values of samples (called particles) to incorporate information from noisy observations of the state of a dynamic model. While the particle filter is a very versatile method it may require a very large number of samples to approximate accurately the conditional density of the state of the model. This has led to considerable research (see e.g. [7, 22]) into how one can modify a particle filter to make it more efficient (see also [3] for a different approach to particle filtering). As an application of the drift relaxation algorithm we show in Section 2 how it can be used to construct a more efficient particle filter.

The paper is organized as follows. Section 1 presents the drift relaxation algorithm for an SDE conditional path sampling problem. Section 2 shows how to use the algorithm to modify a particle filter. Section 3 contains numerical results for the application of the modified particle filter to the standard example of filtering a diffusion in a double-well potential (more elaborate examples will be presented in [13]). Finally, Section 4 discusses

the results as well as current and future work.

1 Conditional path sampling and drift relaxation

Suppose that we are given a system of stochastic differential equations (SDEs)

$$dX_t = a(X_t)dt + \sigma(X_t)dB_t, \quad (1)$$

Suppose also that we want to construct, in the time interval $[0, T]$, sample paths from (1) such that the endpoints are distributed according to the densities $h(X_0)$ and $g(X_T)$ respectively. Equation (1) can be discretized in the interval $[0, T]$ by some numerical approximation scheme [9]. Suppose that we have discretized the interval $[0, T]$ using a stepsize $\Delta t = T/I$. To construct conditional paths of (1) we can start by sampling initial conditions from the density h . For each initial condition X_0 we can then produce the desired conditional paths by sampling the density

$$\prod_{i=1}^I p(X_{T_i}|X_{T_{i-1}})g(X_T), \quad (2)$$

where $p(X_{T_i}|X_{T_{i-1}})$ is the transition probability from $X_{T_{i-1}}$ at time T_{i-1} to the point X_{T_i} at time T_i . The density (2) can be sampled using MCMC assuming that the transition densities $p(X_{T_i}|X_{T_{i-1}})$ can be evaluated. However, the major issue with the MCMC sampling is whether it can be performed efficiently (see e.g. [21, 3]). Instead of MCMC sampling directly from the density (2) i.e., starting from an arbitrary initial path and modifying it to become a path corresponding to (2), we can aid the MCMC sampling process by providing the MCMC sampler of the density (2) with a better initial condition.

To this end, consider an SDE system with modified drift

$$dY_t = b(Y_t)dt + \sigma(Y_t)dB_t, \quad (3)$$

where $b(Y_t)$ can be suitably chosen to facilitate the conditional path sampling problem.

Also, consider the collection of $L + 1$ modified SDE systems

$$dY_t^l = (1 - \epsilon_l)b(Y_t^l)dt + \epsilon_l a(Y_t^l)dt + \sigma(Y_t^l)dB_t,$$

where $\epsilon_l \in [0, 1]$, $l = 0, \dots, L$, with $\epsilon_l < \epsilon_{l+1}$, $\epsilon_0 = 0$ and $\epsilon_L = 1$. Instead of sampling directly a (conditional) path from the SDE (1), one can sample a

path from the modified SDE (3) and gradually morph the path into a path of (1).

Drift relaxation algorithm:

- ($l = 0$) Begin with a sample path from the modified SDE (3).
- Sample through MCMC the density $\prod_{i=1}^I p(Y_{T_i}^0 | Y_{T_{i-1}}^0) g(Y_T^0)$.
- For $l = 1, \dots, L$ take the last sample path from the $(l - 1)$ st SDE and use it as in initial condition for MCMC sampling the density

$$\prod_{i=1}^I p(Y_{T_i}^l | Y_{T_{i-1}}^l) g(Y_T^l)$$

at the l th level.

- Keep the last sample path at the L th level.

The drift relaxation algorithm is similar to Simulated Annealing (SA) used in equilibrium statistical mechanics [10]. However, instead of modifying a temperature as in SA, here we modify the drift of the system.

2 Application to particle filtering

We show in this section how the drift relaxation algorithm can be applied to particle filtering with the aim of bringing the samples closer to the observations.

2.1 Generic particle filter

Suppose that we are given an SDE system and that we also have access to noisy observations Z_{T_1}, \dots, Z_{T_K} of the state of the system at specified instants T_1, \dots, T_K . The observations are functions of the state of the system, say given by $Z_{T_k} = G(X_{T_k}, \xi_k)$, where $\xi_k, k = 1, \dots, K$ are mutually independent random variables. For simplicity, let us assume that the distribution of the observations admits a density $g(X_{T_k}, Z_{T_k})$, i.e., $p(Z_{T_k} | X_{T_k}) \propto g(X_{T_k}, Z_{T_k})$.

The filtering problem consists of computing estimates of the conditional expectation $E[f(X_{T_k}) | \{Z_{T_j}\}_{j=1}^k]$, i.e., the conditional expectation of the state of the system given the (noisy) observations. Equivalently, we are looking to compute the conditional density of the state of the system given

the observations $p(X_{T_k}|\{Z_{T_j}\}_{j=1}^k)$. There are several ways to compute this conditional density and the associated conditional expectation but for practical applications they are rather expensive.

Particle filters fall in the category of importance sampling methods. Because computing averages with respect to the conditional density involves the sampling of the conditional density which can be difficult, importance sampling methods proceed by sampling a reference density $q(X_{T_k}|\{Z_{T_j}\}_{j=1}^k)$ which can be easily sampled and then compute the weighted sample mean

$$E[f(X_{T_k})|\{Z_{T_j}\}_{j=1}^k] \approx \frac{1}{N} \sum_{n=1}^N f(X_{T_k}^n) \frac{p(X_{T_k}^n|\{Z_{T_j}\}_{j=1}^k)}{q(X_{T_k}^n|\{Z_{T_j}\}_{j=1}^k)}$$

or the related estimate

$$E[f(X_{T_k})|\{Z_{T_j}\}_{j=1}^k] \approx \frac{\sum_{n=1}^N f(X_{T_k}^n) \frac{p(X_{T_k}^n|\{Z_{T_j}\}_{j=1}^k)}{q(X_{T_k}^n|\{Z_{T_j}\}_{j=1}^k)}}{\sum_{n=1}^N \frac{p(X_{T_k}^n|\{Z_{T_j}\}_{j=1}^k)}{q(X_{T_k}^n|\{Z_{T_j}\}_{j=1}^k)}}, \quad (4)$$

where N has been replaced by the approximation

$$N \approx \sum_{n=1}^N \frac{p(X_{T_k}^n|\{Z_{T_j}\}_{j=1}^k)}{q(X_{T_k}^n|\{Z_{T_j}\}_{j=1}^k)}.$$

Particle filtering is a recursive implementation of the importance sampling approach. It is based on the recursion

$$p(X_{T_k}|\{Z_{T_j}\}_{j=1}^k) \propto g(X_{T_k}, Z_{T_k}) p(X_{T_k}|\{Z_{T_j}\}_{j=1}^{k-1}), \quad (5)$$

$$\text{where } p(X_{T_k}|\{Z_{T_j}\}_{j=1}^{k-1}) = \int p(X_{T_k}|X_{T_{k-1}}) p(X_{T_{k-1}}|\{Z_{T_j}\}_{j=1}^{k-1}) dX_{T_{k-1}}. \quad (6)$$

If we set

$$q(X_{T_k}|\{Z_{T_j}\}_{j=1}^k) = p(X_{T_k}|\{Z_{T_j}\}_{j=1}^{k-1}),$$

then from (5) we get

$$\frac{p(X_{T_k}|\{Z_{T_j}\}_{j=1}^k)}{q(X_{T_k}|\{Z_{T_j}\}_{j=1}^k)} \propto g(X_{T_k}, Z_{T_k}).$$

The approximation in expression (4) becomes

$$E[f(X_{T_i})|\{Z_{T_j}\}_{j=1}^k] \approx \frac{\sum_{n=1}^N f(X_{T_k}^n) g(X_{T_k}^n, Z_{T_k})}{\sum_{n=1}^N g(X_{T_k}^n, Z_{T_k})} \quad (7)$$

From (7) we see that if we can construct samples from the predictive distribution $p(X_{T_k}|\{Z_{T_j}\}_{j=1}^{k-1})$ then we can define the (normalized) weights $W_{T_k}^n = \frac{g(X_{T_k}^n, Z_{T_k})}{\sum_{n=1}^N g(X_{T_k}^n, Z_{T_k})}$, use them to weigh the samples and the weighted samples will be distributed according to the posterior distribution $p(X_{T_k}|\{Z_{T_j}\}_{j=1}^k)$.

In many applications, most samples will have a negligible weight with respect to the observation, so carrying them along does not contribute significantly to the conditional expectation estimate (this is the problem of degeneracy [10]). To create larger diversity one can resample the weights to create more copies of the samples with significant weights. The particle filter with resampling is summarized in the following algorithm due to Gordon *et al.* [8].

Particle filter

1. Begin with N unweighted samples $X_{T_{k-1}}^n$ from $p(X_{T_{k-1}}|\{Z_{T_j}\}_{j=1}^{k-1})$.
2. **Prediction:** Generate N samples $X_{T_k}^m$ from $p(X_{T_k}|X_{T_{k-1}})$.
3. **Update:** Evaluate the weights

$$W_{T_k}^n = \frac{g(X_{T_k}^m, Z_{T_k})}{\sum_{n=1}^N g(X_{T_k}^m, Z_{T_k})}.$$

4. **Resampling:** Generate N independent uniform random variables $\{\theta^n\}_{n=1}^N$ in $(0, 1)$. For $n = 1, \dots, N$ let $X_{T_k}^n = X_{T_k}^{t_j}$ where

$$\sum_{l=1}^{j-1} W_{T_k}^l \leq \theta^j < \sum_{l=1}^j W_{T_k}^l$$

where j can range from 1 to N .

5. Set $k = k + 1$ and proceed to Step 1.

The particle filter algorithm is easy to implement and adapt for different problems since the only part of the algorithm that depends on the specific dynamics of the problem is the prediction step. This has led to the particle filter algorithm's increased popularity [5]. However, even with the resampling step, the particle filter can still need a lot of samples in order to describe accurately the conditional density $p(X_{T_k}|\{Z_{T_j}\}_{j=1}^k)$. Snyder *et al.* [16] have shown how the particle filter can fail in simple high dimensional

problems because one sample dominates the weight distribution. The rest of the samples are not in statistically significant regions. Even worse, as we will show in the numerical results section, there are simple examples where not even one sample is in a statistically significant region. In the next subsection we present how drift relaxation can be used to push samples closer to statistically significant regions.

2.2 Particle filter with MCMC step

Several authors (see e.g. [7, 22]) have suggested the use of a MCMC step after the resampling step (Step 4) in order to move samples away from statistically insignificant regions. There are many possible ways to append an MCMC step after the resampling step in order to achieve that objective. The important point is that the MCMC step must preserve the conditional density $p(X_{T_k} | \{Z_{T_j}\}_{j=1}^k)$. In the current section we show that the MCMC step constitutes a case of conditional path sampling.

We begin by noting that one can use the resampling step (Step 4) in the particle filter algorithm to create more copies not only of the good samples according to the observation, but also of the values (initial conditions) of the samples at the previous observation. These values are the ones who have evolved into good samples for the current observation (see more details in [22]). The motivation behind producing more copies of the pairs of initial and final conditions is to use the good initial conditions as starting points to produce statistically more significant samples according to the current observation. This process can be accomplished in two steps. First, Step 4 of the particle filter algorithm is replaced by

Resampling: Generate N independent uniform random variables $\{\theta^n\}_{n=1}^N$ in $(0, 1)$. For $n = 1, \dots, N$ let $(X_{T_{k-1}}^n, X_{T_k}^n) = (X_{T_{k-1}}'^j, X_{T_k}'^j)$ where

$$\sum_{l=1}^{j-1} W_{T_k}^l \leq \theta^j < \sum_{l=1}^j W_{T_k}^l$$

Also, through Bayes' rule [22] one can show that the posterior density $p(X_{T_k} | \{Z_{T_j}\}_{j=1}^k)$ is preserved if one samples from the density

$$g(X_{T_k}, Z_{T_k})p(X_{T_k} | X_{T_{k-1}})$$

where $X_{T_{k-1}}$ are given by the modified resampling step. This is a problem of conditional path sampling for (continuous-time or discrete) stochastic

systems. The important issue is to perform the necessary sampling efficiently [3, 22]. We propose to do that here using drift relaxation (see Section 1). The particle filter with MCMC step algorithm is given by

Particle filter with MCMC step

1. Begin with N unweighted samples $X_{T_{k-1}}^n$ from $p(X_{T_{k-1}}|\{Z_{T_j}\}_{j=1}^{k-1})$.
2. **Prediction:** Generate N samples $X_{T_k}^n$ from $p(X_{T_k}|X_{T_{k-1}})$.
3. **Update:** Evaluate the weights

$$W_{T_k}^n = \frac{g(X_{T_k}^n, Z_{T_k})}{\sum_{n=1}^N g(X_{T_k}^n, Z_{T_k})}.$$

4. **Resampling:** Generate N independent uniform random variables $\{\theta^n\}_{n=1}^N$ in $(0, 1)$. For $n = 1, \dots, N$ let $(X_{T_{k-1}}^n, X_{T_k}^n) = (X_{T_{k-1}}^{l_j}, X_{T_k}^{l_j})$ where

$$\sum_{l=1}^{j-1} W_{T_k}^l \leq \theta^j < \sum_{l=1}^j W_{T_k}^l$$

where j can range from 1 to N .

5. **MCMC step:** For $n = 1, \dots, N$ choose a modified drift (possibly different for each n). Construct a path for the SDE with the modified drift starting from $X_{T_{k-1}}^n$. Construct through drift relaxation a Markov chain for $Y_{T_k}^n$ with stationary distribution

$$g(Y^n, Z_{T_k})p(Y^n|X_{T_{k-1}}^n)$$

6. Set $X_{T_k}^n = Y_{T_k}^n$.
7. Set $k = k + 1$ and proceed to Step 1.

Since the samples $X_{T_k}^n = Y_{T_k}^n$ are constructed by starting from different sample paths, they are independent. Also, note that the samples $X_{T_k}^n$ are unweighted. However, we can still measure how well these samples approximate the posterior density by comparing the effective sample sizes of the particle filter with and without the MCMC step. For a collection of N samples the effective sample size $ess(T_k)$ is defined by

$$ess(T_k) = \frac{N}{1 + C_k^2}$$

where

$$C_k = \frac{1}{W_k} \sqrt{\frac{1}{N} \sum_{n=1}^N (g(X_{T_k}^n, Z_{T_k}) - W_k)^2} \text{ and } W_k = \frac{1}{N} \sum_{n=1}^N g(X_{T_k}^n, Z_{T_k}).$$

The effective sample size can be interpreted as that the N weighted samples are worth of $ess(T_k) = \frac{N}{1+C_k^2}$ i.i.d. samples drawn from the target density, which in our case is the posterior density. By definition, $ess(T_k) \leq N$. If the samples have uniform weights, then $ess(T_k) = N$. On the other hand, if all samples but one have zero weights, then $ess(T_k) = 1$.

3 Numerical results

We present numerical results of the particle filter algorithm with MCMC step for the standard problem of diffusion in a double-well potential (more elaborate applications of the method will be presented elsewhere [13]). Our objective here is to show how the generic particle filter's performance can be significantly improved by incorporating the MCMC step via drift relaxation.

The problem of diffusion in a double well potential is described by the scalar SDE

$$dX_t = -4X_t(X_t^2 - 1) + \frac{1}{2}dB_t. \quad (8)$$

The deterministic part (drift) describes a gradient flow for the potential $U(x) = x^4 - 2x^2$ which has two minima, at $x = \pm 1$. In the notation of Section 2 we have $a(X_t) = -4X_t(X_t^2 - 1)$ and $\sigma(X_t) = \frac{1}{2}$. If the stochastic term is zero the solution wanders around one of the minima depending on the value of the initial condition. A weak stochastic term leads to rare transitions between the minima of the potential. We have chosen the coefficient $\frac{1}{2}$ to make the stochastic term rather weak. This is done because we plan to enforce the observations to alternate among the minima, and thus check if the particle filter can track these transitions.

The SDE (8) is discretized by the Euler-Maruyama [9] scheme with step size $\Delta t = 10^{-2}$ which is small enough to guarantee stability of the scheme. The initial condition is set to -1 and there is a total of 10 observations at $T_k = k, k = 1, \dots, 10$. The observations are given by $Z_{T_k} = X_{T_k} + \xi_k$, where $\xi_k \sim N(0, 0.01)$ for $k = 1, \dots, 10$. For this choice of observation noise, the observation density (also called likelihood) is given by

$$g(X_{T_k}, Z_{T_k}) \propto \exp \left[-\frac{(Z_{T_k} - X_{T_k})^2}{2 * 0.01} \right] \quad (9)$$

The observations alternate between 1 and -1 . In particular, for $k = 1, \dots, 10$ we have $Z_{T_k} = -1$ if k is odd and $Z_{T_k} = 1$ if k is even.

In order to apply the MCMC step with drift relaxation we need to define the modified drift $b(Y_t)$ for the process Y_t given by

$$dY_t = b(Y_t) + \frac{1}{2}dB_t. \quad (10)$$

The modified drift can be the same for all the samples or different for each sample. Since the difficulty in tracking the observations comes from the inability of the original SDE (8) to make frequent transitions between the two minima of the double well, an intuitively appealing choice for $b(Y_t)$ is $b(Y_t) = -\alpha 4Y_t(Y_t^2 - 1)$, where $\alpha < 1$. This drift corresponds to the potential $W(y) = \alpha(y^4 - 2y^2)$. The potential $W(y)$ has its minima also located at $y = \pm 1$. However, the value of the potential at the minima is $-\alpha$ instead of -1 for the potential $U(x)$. This means that the wells corresponding to the minima of $W(y)$ are shallower than the wells corresponding to the minima of $U(x)$. This makes the transitions between the two wells for the process Y_t more frequent than for the original process X_t . For the numerical experiments we have chosen $\alpha = 0.1$.

The sequence of modified SDEs for the drift relaxation algorithm with L levels is given by

$$dY_t^l = (1 - \epsilon_l)b(Y_t^l)dt + \epsilon_l a(Y_t^l)dt + \frac{1}{2}dB_t, \quad (11)$$

where $\epsilon_l \in [0, 1]$, $l = 0, \dots, L$, with $\epsilon_l < \epsilon_{l+1}$, $\epsilon_0 = 0$ and $\epsilon_L = 1$. For our numerical experiments we chose $L = 10$ and $\epsilon_l = l/10$.

Recall that the density we want to sample during the MCMC step is given by

$$g(X_{T_k}, Z_{T_k})p(X_{T_k}|X_{T_{k-1}}),$$

where $p(X_{T_k}|X_{T_{k-1}})$ is the transition probability between $X_{T_{k-1}}$ and X_{T_k} . For many applications, sampling directly from $p(X_{T_k}|X_{T_{k-1}})$ may be impossible. Thus, one needs to resort to some numerical approximation scheme which approximates the path between $X_{T_{k-1}}$ and X_{T_k} by a discretized path. However (see [22] for details), even the evaluation of the discretized path's density may not be efficient. Instead, by using the fact that each Brownian path in (8) gives rise to a unique path for X_t [15], we can replace the

sampling of $g(X_{T_k}, Z_{T_k})p(X_{T_k}|X_{T_{k-1}})$ by sampling from the density

$$\exp\left[-\frac{(Z_T - X_T^n(\{\Delta B_i^n\}_{i=0}^{I-1}))^2}{2 * 0.01}\right] \prod_{i=0}^{I-1} \exp\left[-\frac{(\Delta B_i^n)^2}{2 * \Delta t}\right] = \exp\left[-\left(\frac{(Z_T - X_T^n(\{\Delta B_i^n\}_{i=0}^{I-1}))^2}{2 * 0.01} + \sum_{i=0}^{I-1} \frac{(\Delta B_i^n)^2}{2 * \Delta t}\right)\right] \quad (12)$$

where $\{\Delta B_i^n\}_{i=0}^{I-1}$ are the Brownian increments of the discretized path connecting $X_{T_{k-1}}$ and X_{T_k} . Also, note that the final point X_{T_k} has now become a function of the entire Brownian path $\{\Delta B_i^n\}_{i=0}^{I-1}$. For the numerical experiments we have chosen $\Delta t = \frac{T_k - T_{k-1}}{I} = 10^{-2}$ which, since $T_k - T_{k-1} = 1$, gives $I = 100$.

We use drift relaxation to produce samples from the density (12). The Markov chain at each level of the drift relaxation algorithm is constructed using Hybrid Monte Carlo (HMC) [10]. At the l th level, we can discretize (11), say with the Euler-Maruyama scheme, and the points on the path will be given by

$$Y_{i\Delta t}^{l,n} = Y_{(i-1)\Delta t}^{l,n} + (1 - \epsilon_l)b(Y_{(i-1)\Delta t}^{l,n})\Delta t + \epsilon_l a(Y_{(i-1)\Delta t}^{l,n})\Delta t + \frac{1}{2}\Delta B_{i-1}^{l,n},$$

for $i = 1, \dots, I$. We can use more sophisticated schemes than the Euler-Maruyama scheme for the discretization of the simplified SDE (10) at the cost of making the expression for the density more complicated.

We can define a potential $V_{\epsilon_l}(\{\Delta B_i^{l,n}\}_{i=0}^{I-1})$ for the variables $\{\Delta B_i^{l,n}\}_{i=0}^{I-1}$. The potential is given by

$$V_{\epsilon_l}(\{\{\Delta B_i^{l,n}\}_{i=0}^{I-1}\}) = \frac{(Z_T - Y_{I\Delta t}^{l,n}(\{\Delta B_i^{l,n}\}_{i=0}^{I-1}))^2}{2 * 0.01} + \sum_{i=0}^{I-1} \frac{(\Delta B_i^{l,n})^2}{2 * \Delta t}$$

and the density to be sampled can be written as

$$\exp\left[-V_{\epsilon_l}(\{\Delta B_i^{l,n}\}_{i=0}^{I-1})\right].$$

The subscript ϵ_l is to denote the dependence of the potential on the drift relaxation parameter ϵ_l . In HMC one considers the variables on which the potential depends as the position variables of a Hamiltonian system. In our case we have I position variables so we can define a I -dimensional position vector $\{q_i\}_{i=1}^I$. The next step is to augment the position variables vector by

a vector of associated momenta $\{p_i\}_{i=1}^I$. Together they form a Hamiltonian system with Hamiltonian given by

$$H_{\epsilon_l}(\{q_i\}_{i=1}^I, \{p_i\}_{i=1}^I) = V_{\epsilon_l}(\{q_i\}_{i=1}^I) + \frac{p^T p}{2},$$

where $p = (p_1, \dots, p_I)$ is the momenta vector. Thus, the momenta variables are Gaussian distributed random variables with mean zero and variance 1. The equations of motion for this Hamiltonian system are given by Hamilton's equations

$$\frac{dq_i}{d\tau} = \frac{\partial H_{\epsilon_l}}{\partial p_i} \text{ and } \frac{dp_i}{d\tau} = -\frac{\partial H_{\epsilon_l}}{\partial q_i} \text{ for } i = 1, \dots, I.$$

Note that the Hamiltonian depends also on the initial condition for each sample Y_0^n and we have written an equation for the evolution of $q_1 = Y_0^n$ as well its associated momentum p_1 . Since the Y_0^n are fixed by the resampling procedure we do not evolve them. However, note that the Brownian increment ΔB_0 needs to be evolved because it affects the evolution of $Y_{\Delta t}$.

HMC proceeds by assigning initial conditions to the momenta variables (through sampling from $\exp(-\frac{p^T p}{2})$), evolving the Hamiltonian system in fictitious time τ for a given number of steps of size $\delta\tau$ and then using the solution of the system to perform a Metropolis accept/reject step (more details in [10]). After the Metropolis step, the momenta values are discarded. The most popular method for solving the Hamiltonian system, which is the one we also used, is the Verlet leapfrog scheme. In our numerical implementation, we did not attempt to optimize the performance of the HMC algorithm. For the sampling at each level of the drift relaxation process we used 10 Metropolis accept/reject steps and 1 HMC step of size $\delta\tau = 10^{-2}$ to construct a trial path. A detailed study of the drift relaxation/HMC algorithm for conditional path sampling problems outside of the context of particle filtering will be presented in a future publication.

For the chosen values of the parameters for the drift relaxation and HMC steps, the particle filter with MCMC step is about 500 times more expensive per sample (particle) than the generic particle filter. However, we show that this increase in cost per sample is worthwhile. Figure 1 compares the performance of the particle filter with MCMC step with 10 samples and the generic particle filter with 5000 samples. It is obvious that the particle filter with MCMC step follows accurately all the transitions between the two minima of the double-well. On the other hand, the generic particle filter captures accurately only every other observation. It fails to performs

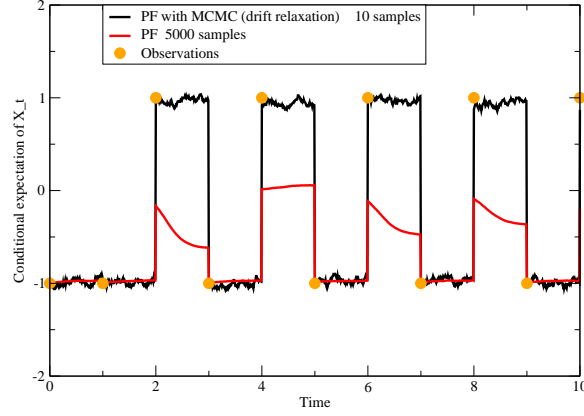


Figure 1: Comparison of the conditional expectation of X_t as computed by the generic particle filter and the particle filter with MCMC step. Different drift for each sample.

the transitions between the two minima of the double-well. Of course, since the particle filter with MCMC step uses only 10 samples the conditional expectation estimate of the hidden signal is not as smooth as the estimate of the generic particle filter which uses 5000 samples. However, the particle filter with MCMC step allows a much better resolution of the conditional density (conditioned on the observations). This can be seen by computing the effective sample size for the two filters. Figure 2 shows the effective sample size for both filters. Because of the different number of samples used in the two filters we have plotted the effective sample size for each filter as a percentage of the respective sample size. We see that the particle filter with MCMC step has overall a much better sample size than the generic particle filter. The generic particle filter has a wildly fluctuating effective sample size. In particular, since the generic particle filter misses every other observation, the corresponding effective sample size dips down to 1 sample for the missed observations. Note from Figure 1 that even this one sample which dominates the observation weight does not come close to the observation. For the observations that the generic particle filter does capture, its effective sample size is still lower than the effective sample size of the particle filter with MCMC step.

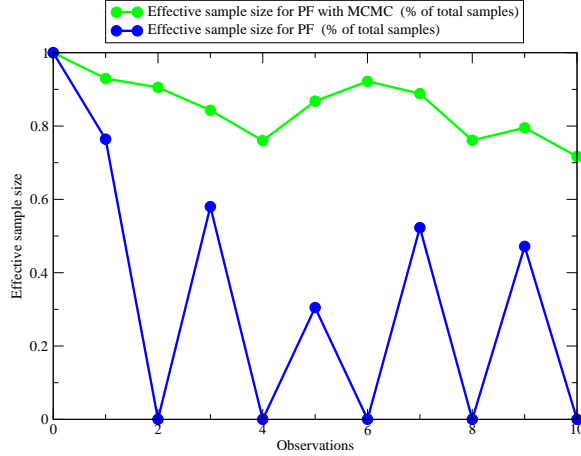


Figure 2: Comparison of the effective sample size as computed by the generic particle filter and the particle filter with MCMC step. A different drift is used for each sample.

4 Discussion

We have presented an algorithm for conditional path sampling of SDEs. The proposed algorithm is based on drift relaxation which allows to sample conditional paths from a modified drift equation. The conditional paths of the modified drift equation are then morphed into conditional paths of the original equation. We have called this process of gradually enforcing the drift of the original equation drift relaxation. The algorithm has been used to create a modified particle filter for SDEs. We have shown that the modified particle filter’s performance is significantly better than the performance of a generic particle filter.

In the current work, we have examined the application of drift relaxation to the filtering problem of diffusion in a double-well potential which is a standard example in the filtering literature. The same algorithm can be applied to the problem of tracking a single target. A problem of great practical interest is that of tracking not only one but multiple moving targets [11, 14, 19, 20]. The multi-target tracking problem is much more difficult than the single-target problem due to the combinatorial explosion of the number of possible target-observation association arrangements. In this context, the accurate tracking of each target becomes crucial. Suppose that only one of the targets is of interest and the rest act as decoys [12]. The inability to track each potential target accurately can lead to ambiguity about

the targets' movement if the observations for different targets are close. We have already applied the drift relaxation modified particle filter to multi-target tracking problems with very encouraging results which will appear elsewhere [13].

Acknowledgements

I am grateful to Profs. V. Maroulas and J. Weare for many helpful discussions and comments. I would also like to thank for its hospitality the Institute for Mathematics and its Applications (IMA) at the University of Minnesota where the current work was completed.

References

- [1] Bolhuis P. G., Chandler D., Dellago C. and Geissler P. L., Transition path sampling: throwing ropes over rough mountain passes in the dark, *Ann. Rev. Phys. Chem.* 53 (2002) pp. 291-318.
- [2] Chorin A.J., Monte Carlo without chains, *Comm. Appl. Math. Comp. Sc.* 3 (2008) pp. 77-93.
- [3] Chorin, A.J. and Tu X., Implicit sampling for particle filters, *Proc. Nat. Acad. Sc. USA* 106 (2009) pp. 17249-17254.
- [4] Deuffhard P. and Bornemann F., *Scientific Computing with Ordinary Differential Equations*, Springer NY, 2002.
- [5] Doucet A., de Freitas N. and Gordon N. (eds.) , *Sequential Monte Carlo Methods in Practice*, Springer NY, 2001.
- [6] Duane S., Kennedy A., Pendleton B. and Roweth D., Hybrid Monte Carlo, *Phys. Lett. B* 195 (1987) pp. 216-222.
- [7] Gilks W. and Berzuini C., Following a moving target - Monte Carlo inference for dynamic Bayesian models, *J. Royal Stat. Soc. B* 63 (1) (1999) pp. 2124-2137.
- [8] Gordon N.J., Salmond D.J. and Smith A.F.M., Novel approach to nonlinear/non-Gaussian Bayesian state estimation, *Proc. Inst. Elect. Eng. F* 140(2) (1993) pp. 107-113.

- [9] Kloeden P. and Platen E., Numerical Solution of Stochastic Differential Equations, Springer Berlin, 1992.
- [10] Liu J.S., Monte Carlo Strategies in Scientific Computing, Springer NY, 2001.
- [11] Mahler R. P. S., Statistical Multisource-Multitarget Information Fusion, Artech House Publishers MA, 2007.
- [12] Mahler R. P. S. and Maroulas V., Tracking spawning objects, (2010) *submitted*.
- [13] Maroulas V. and Stinis P., A drift relaxation Monte Carlo approach to particle filtering for multi-target tracking, (2010), arXiv:1006.3100v1.
- [14] Ng W., Li J.F., Godsill S.J. and Vermaak J., A hybrid approach for online joint detection and tracking for multiple targets, Proc. IEEE Aerospace Conference (2005) pp. 2126-2141.
- [15] Oksendal B., Stochastic Differential Equations, An Introduction with Applications, Sixth Edition Springer Heidelberg, 2005.
- [16] Snyder C., Bengtsson T., Bickel P. and Anderson J., Obstacles to High-dimensional Particle Filtering, Mon. Wea. Rev., Vol. 136 (2008) pp. 4629-4640.
- [17] Stoer J. and Bulirsch R., Introduction to Numerical Analysis, Third Edition, Springer 2002.
- [18] Stuart A.M., Voss J. and Wiberg P., Conditional path sampling of SDEs and the Langevin MCMC method, Comm. Math. Sci., Vol.2 No.4 (2004) pp. 685-697.
- [19] Vermaak J., Godsill S. and Perez P., Monte Carlo filtering for multi-target tracking and data association, IEEE Trans. Aero. Elect. Sys., 41(1) (2005) pp. 309-332.
- [20] Vo B-N., Singh S. and Doucet A., Sequential Monte Carlo Methods for Multi-Target Filtering with Random Finite Sets, IEEE Trans. Aero. Elect. Sys., 41(4) (2005) pp. 1224-1245.
- [21] Weare J., Efficient Monte Carlo sampling by parallel marginalization, Proc. Natl. Acad. Sci. USA 104(31) (2007) pp. 12657-12662

- [22] Weare J., Particle filtering with path sampling and an application to a bimodal ocean current model, J. Comp. Phys. 228 (2009) pp. 4312-4331.